AD A051 320

Semiannual Technical Summary

Information Processing Techniques Program

Volume II:

Communications-Adaptive Internetting

BEST AVAILABLE COPY

30 September 1977

Prepared for the Defense Advanced Research Projects Agency
under Electronic Systems Division Contract F19628-76-C-0002 by

Lincoln Laboratory

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LEXINGTON, MASSACHUSETTS

[illegible] Lincoln Laboratory, [illegible] Institute of Technology. [illegible] the Defense Advanced Research Projects Agency [illegible] Force Contract F19628-[illegible]-C-0002 (ARPA Order [illegible]).

[illegible] may be reproduced to satisfy needs of U.S. Government agencies.

The views and conclusions contained in this document are those of the contractor and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER

Raymond L. Loiselle

Raymond L. Loiselle, Lt. Col., USAF
Chief, ESD Lincoln Laboratory Project Office

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LINCOLN LABORATORY

INFORMATION PROCESSING TECHNIQUES PROGRAM
VOLUME II: COMMUNICATIONS-ADAPTIVE INTERNETTING

SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1 APRIL – 30 SEPTEMBER 1977

ISSUED 19 JANUARY 1978

Approved for public release; distribution unlimited.

LEXINGTON MASSACHUSETTS

ABSTRACT

This report describes work performed on the Communications-Adaptive Internetting program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 April through 30 September 1977.

CONTENTS

COMMUNICATIONS-ADAPTIVE INTERNETTING

I. INTRODUCTION AND SUMMARY

The goal of this program is investigation and development of techniques for communications-adaptive internetting with particular emphasis on digital voice communications. The specific class of problems addressed relates to packet speech networks whose interconnecting links may be stressed as a result of traffic overloads or are time varying due to natural or hostile actions. This program extends the technology of fixed-topology packet-switching speech communications networks by introducing techniques for adapting both source encoder algorithms and network utilization strategies to the time-varying character of the links in an attempt to provide sustained voice communications capability under adverse conditions.

The major accomplishments of the past six months are summarized below.

A. Demonstration System for Packetized Voice over a Single-User Time-Varying Link

The system is operational and supports variable-rate speech. The real-time variable-rate modem implementation on the Lincoln Digital Voice Terminal (LDVT) is not yet complete and, therefore, simpler LDVT programs which simulate the effects of these modems have been included temporarily in the demonstration system. The setup of input/output communication between the PDP-11/45 and the modems is different from that described in the previous SATS,* resulting in some hardware savings but substantial software changes. In order to make the system work under the UNIX operating system, substantial buffers (and resulting delays) between transmitter and receiver had to be introduced. These delays have caused the demonstration system to exhibit slower response to capacity changes than might be anticipated in more tightly controlled real environments.

B. Adaptive Voice Communications in a Wideband Multi-user Environment

A variable-rate speech system without a simulated modem has been combined with a real-time PDP-11 program capable of discarding packets or introducing bit errors based on previously recorded measurements, which can be obtained through non-real-time simulation of a multi-user network. As an example application of this technique, an existing simulation of a Packetized Virtual Circuit (PVC) network was modified to provide packet loss history for a single user, and this record was applied to the real-time simulation.

In preparation for a planned experiment involving adaptive speech communication in a broadcast satellite environment, a preliminary analysis of the required parameters and costs of satellite ground stations for that experiment has been carried out.

* Information Processing Techniques Program Semiannual Technical Summary, Volume II: Communications-Adaptive Internetting, Lincoln Laboratory, M.I.T. (31 March 1977), DDC AD-A044071.

II. DEMONSTRATION SYSTEM FOR PACKETIZED VOICE OVER A SINGLE-USER TIME-VARYING LINK

A. Final System Configuration

The demonstration system is operational and supports real-time speech. The final system configuration shown in Fig. 1 differs somewhat from the configuration described in the previous SATS in that only one (instead of two) DR11B is used to interconnect the PDP-11 with the modem chain, and additional LDVTs are used as buffers for the modulator and demodulator. This change has not affected system performance. The flexible signal conditioner has been completed and incorporated into the system to facilitate vocoder bit rate changes. The real-time variable-rate modulator and demodulator implementations are not quite complete, and temporarily have been replaced by simpler programs which simulate the effect of the modem.

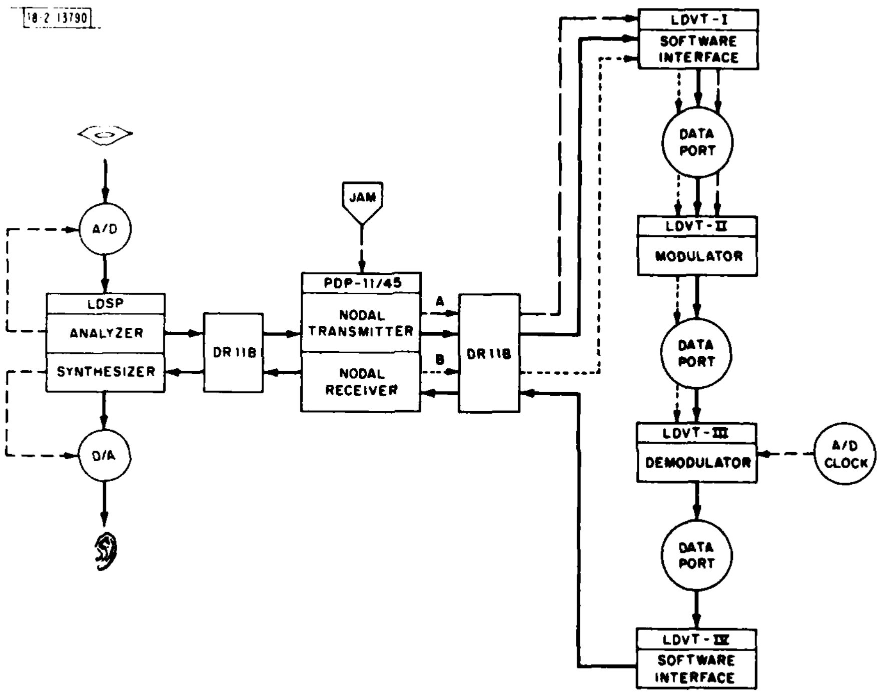

Fig. 1. Half-duplex variable-rate speech system.

Referring to Fig. 1, the new system configuration involved a major change in the format of information to the modem from the nodal processor and the inclusion of software interfaces in LDVT-I and LDVT-IV to accommodate the existence of only one DR11B instead of two between the PDP-11/45 and the simulated modem. Information from the demodulator to the nodal processor remains basically unchanged. However, information to the modulator now consists of blocks of 128 symbols packed four to a word preceded by two words of control information. The first control word (transmitted along the path marked A in Fig. 1) contains rate and noise level for the modulator and the second (transmitted along the path marked B in Fig. 1) contains the independently determined rate for the demodulator. The latter will be passed periodically from

the modulator to the demodulator via a data port as indicated. Because changes in transmission rate occur only at packet boundaries, the inclusion of the transmission rate only once per half packet does not affect system performance.

The function of the software interfaces in LDVT-I and LDVT-IV is to provide for the modem a continuous source of bits and an unlimited buffer for bits, which was the basic requirement for a modem whose transmitter and receiver will be operating at different rates during responses to changes in channel capacity by the nodal processor. Thus, the absence of a second DR11B becomes transparent to the modem, and the system is able to perform as originally intended. The A/D clock attached to LDVT-III serves as the master clock for the entire system, determining the basic unit of data transfer from the PDP-11/45 to the LDVTs and back to the PDP-11/45.

B. Management of Input/Output under UNIX

The UNIX time-sharing system which supports this demonstration prohibits the user from directly issuing any I/O commands. The system software (I/O driver) which performs these functions maintains offsets to one or two circular tables of pointers (in the user area) to a specified number of fixed length blocks, also in the user area. The user program can dynamically change the pointers in these tables. The driver may be operated in several modes. For example, the "alternate read/write" mode requires two circular tables of pointers, one pointing to input blocks and the other pointing to output blocks. The driver will initiate block transfers alternately to and from the peripheral device. When a block transaction is completed, the corresponding system offset will be updated. This offset is monitored by the user program but cannot be changed by it. Once a block transfer is initiated by the driver, data transfer takes approximately 6 μsec per word between the peripheral device and the DR11B, a direct memory access device which uses central processor unit (CPU) cycles as needed. The system overhead involved in setting up the next block transfer is approximately 1 msec. The DR11B's shown in Fig. 1 communicate with either a Lincoln Digital Signal Processor (LDSP) or two LDVTs through a pair of 64-word FIFOs (first-in/first-out buffers).

Data transfer to and from the vocoder in the LDSP is achieved by alternately writing to and reading from the LDSP (via the FIFOs) blocks of 60 words. Each block represents the data portion of a 64-word packet. While one FIFO is being emptied by the LDSP, the other is being filled. A small amount of buffering is necessary in each vocoder to absorb the total system overhead for the PDP-11 interaction with the LDSP and the LDVTs. Data transfer to or from the PDP-11/45 FIFOs takes place upon request by the LDSP and is essentially instantaneous. Upon the completion of a data transaction with the PDP-11/45, the LDSP receives an interrupt; and an interrupt service routine in the LDSP records this acknowledgment. The LDSP interrupt service routines for the A/D and D/A converters constantly monitor the status of the analysis and synthesis buffers, respectively. If there is a word to send or room to receive a word, the appropriate request is issued to the PDP-11/45 only if the previous input or output request has been acknowledged. This ensures that data transactions will catch up after a FIFO has been unable to accommodate the LDSP during the time the I/O drivers have been initiating new block transfers.

Alternately writing to the modulator and reading from the demodulator is not feasible because their rates are independent, and the symmetry of data transactions existing between the LDSP and PDP-11/45 is not maintained in this data exchange. Even when the rates are equal, the buffering needed to absorb the system overhead between blocks and the bookkeeping necessary

to guarantee readiness of the FIFOs is intolerable due to the amount of computation required to simulate the modem and the inflexibility of the I/O facility in the LDVTs. The I/O driver may be operated in an "asynchronous read/write" mode in which if a block transfer to a FIFO is completed and the alternate FIFO does not require servicing, the original FIFO is checked for servicing before the driver "goes to sleep." Provision is made for prodding the driver from the user program in order to force periodic checks on the status of the FIFOs. Using this mode, however, does not alone solve the problem. It is essential that both FIFOs be ready to service the modem upon request to prevent hanging up the modem processors; it is also mandatory to provide for erratic data exchange which will occur when the modem transmitter and receiver rates differ and the receiver drifts some number of symbols due to resynchronization. To meet these requirements, the following solution was reached. Blocks of data to and from the modem represent one-half of a packet or 512 bits. A block of 34 words to the modulator includes 128 symbols and two control words, one for the modulator and one to be passed to the demodulator as described previously. A block of 64 words from the demodulator contains 128 symbols and the corresponding capacity assessments. The critical ingredient here is that the block sizes are less than or equal to the size of a FIFO. Because the input block is greater than one-half of a FIFO, the following trick is played in the software interface in LDVT-I. The interface begins by reading 38 words all at once from the PDP-11/45 into a buffer from which the first 34 words will be sent to the modulator in LDVT-II upon request by the modulator. When 34 words have been read by the modulator, LDVT-I will read the next 34 words from the PDP-11/45, and so on. Therefore, if the FIFO from the PDP-11/45 has any room, it has room for a full block. Similarly, the interface in LDVT-IV accepts 64 words from the demodulator in LDVT-III upon request by the demodulator, after which the interface sends 64 words at once to the PDP-11/45. Thus, at no time will the driver be "stuck" on a partially serviceable FIFO, and the asymmetric data exchange is accommodated. The simulated modem and interfaces residing in the four LDVTs will hereafter be referred to as the modem cascade.

Because the system offset to a table of block pointers maintained by an I/O driver indicates only which block transaction is in progress and not how many word transactions have actually taken place, a certain amount of buffering is necessary to avoid disaster. In order to minimize degradation of system behavior due to discrepancies in the A/D clock driving the modem and that driving the vocoder, two sets of four buffers are used to buffer data to and from the vocoder. After starting both the I/O driver which services the modem cascade and that which services the vocoder, the nodal transmitter and receiver in the PDP-11/45 independently monitor the system offsets to the modulator and demodulator buffers, respectively. Until a fixed number of block transactions with the modem cascade has occurred, processing of data to or from the modem cascade does not begin. This start-up procedure is a way of synchronizing modem transactions and nodal processing, necessary because of the minimal amount of buffering to and from the modem cascade. This procedure also is used each time a new vocoder has been loaded into the LDSP, and it has the additional function of allowing time for the modem to adjust to the new rate. When the required number of transactions has been satisfied, the nodal processor initializes the corresponding user offset of the I/O driver servicing the vocoder and begins processing data to or from the modem cascade. By initializing the user offset such that if the first vocoder block is active, processing will be done on the third buffer, clock discrepancies would have to result in the drift of one packet's worth of speech data before the nodal processor would

be forced to insert or discard a data block. Independent of modem activity, the system and user offsets to the buffers to and from the vocoder are constantly monitored; and necessary adjustments to the user offsets are made to maintain this cushion of buffering.

C. Speech Delays and Response Time of System to Capacity Changes

The effect of the total amount of buffering in the system is twofold. Not only is the speech delayed, but response time to changes in channel capacity is a pessimistic indication of what could be expected from a tighter system with direct control of I/O activity. In Fig.2, time has

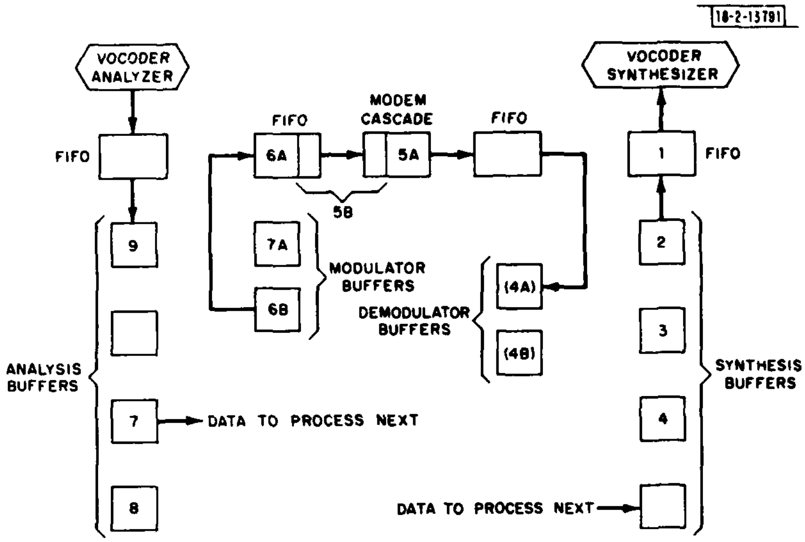

Fig. 2. Packet delay due to buffering of data through the system.

been brought to a standstill in order to illustrate the flow of data through the system and the consequences of the necessary buffering. Assume the following possible status of block activity in Fig.2:

(1) Block 1 (already in the FIFO) is being processed by the vocoder, and block 2 is indicated by the driver as the block in progress.

(2) The second half of packet 4 (4B) has just been received from the modem cascade, and block 4 is now ready to be sent to the vocoder.

(3) The modem cascade is processing the first half of packet 5 (5A) which will not be written to the FIFO until the entire half packet's worth of data have been processed.

(4) Half packets 5B and 6A are already in the FIFO to the modem cascade and await processing.

(5) Half packets 6B and 7A are ready to be sent to the FIFO to the modem cascade.

(6) The remainder of block 7 awaits packetizing upon completion of the next half-packet transaction with the FIFO to the modem cascade.

(7) Block 9 is being read from the FIFO as vocoder data become available.

The nine-packet delay of speech is less than 0.5 sec at the highest rate and greater than 3.5 sec at the lowest rate. Referring to Fig.2, if channel capacity were to increase at this moment, the change in noise level would accompany half packet 7B; however, the new capacity assessment would not be made by the nodal receiver until packet 8 was received, resulting in a response delay four times that to be expected in a tight system. Assuming instantaneous reply on the imaginary return link, the message from the nodal transmitter to increase the channel rate would not reach the nodal receiver until another four packets had passed through the system. Because closing the driver to the vocoder in order to load a new algorithm into the LDSP and reopening the driver to the vocoder interrupts the flow of data to the modem cascade, four packets are allowed to flow through the modem cascade before nodal processing resumes at the new rate, as previously described. Again, a tighter system would provide less delay after the new vocoder was loaded. A decrease in capacity would cause the modulator to distort its signal once it had received the change in noise level. Upon detecting three consecutive packets in error, the nodal receiver would drop to the lowest rate and send a "drop" message to the nodal transmitter via the imaginary return link, causing the nodal transmitter to drop its rate to low. However, the change in both rates would not reach the modem until three packets had trickled through the system. If capacity had dropped to slightly below the original rate of channel operation, the remaining three packets might be received correctly. To prevent the nodal receiver from assuming that the drop message got through and that communication has been established at the low rate, four consecutive blocks must be received correctly before the nodal receiver sends a capacity assessment reply to the nodal transmitter. If the nodal transmitter has received a drop message previously, a reply will force the transmitter to append a rate change message to the next packet even if the new rate remains low. (Be reminded that the nodal receiver constantly sends replies over the imaginary link and that unless the nodal transmitter has received a drop message, a reply from the nodal receiver results in a rate change message only when the reply rate is greater than the rate of channel operation.) A new vocoder is loaded into the LDSP any time the nodal receiver gets a rate change message. Communication at the lowest rate is always assumed to be free of errors.

In summary, the delays described are peculiar to the variable rate system supported by the available PDP-11/45 facility and, in general, will cause longer duration switching transients than would occur in a real system with tight I/O control. More realistic switching delays can be obtained by using the variable-rate speech system without a modem which is described in Sec.III-A. However, the variable-rate system with modem represents a more complete demonstration of adaptive speech communication over a single time-varying link, substantiating the viability of the channel probing and adaptation techniques in the modem (see previous SATS), and addressing in detail the interplay among the changing channel conditions, the rate-adaptive modem, and the variable-rate vocoder.

D. Modem and Vocoder Rates

The A/D clock attached to LDVT-III which drives the modem may be set to an accuracy of 0.5 μsec. A unit of 97.5 μsec per transmission of 16 complex pairs from modulator to demodulator has been chosen to approximate bit rates ranging from 20,480 to 2560 bps. The actual

TABLE I MODEM AND VOCODER RATES			
Modem Rate (bps)	Vocoder Rate (bps)		Vocoder Sampling Interval (μsec)
	Supported	Actual	
20,512.82	19,230.77	19,230.76	52.0
10,256.41	9,615.38	9,615.38	104.0
5,128.21	4,807.69	4,807.41	131.1
2,564.10	2,403.85	2,404.02	131.5

modem rates and the vocoder rates supported are shown in Table I. The vocoder rates have been adjusted as indicated by tweaking the sampling intervals in the LDSP signal conditioner which may be set to an accuracy of 50 nsec. The actual vocoder rates are so close to the supported bit rates that the drift of one packet's worth of data would occur only once every 57 min. in the worst case. Thus, these slight rate discrepancies will have no ill effect on the demonstration.

E. System Refinements

Experience with the system has led to some refinements in both the linear predictive coding (LPC) algorithms and the nodal processor. When a new vocoder has been loaded into the LDSP, if data which have been generated by the previous vocoder are allowed to trickle through the system during the start-up procedure discussed previously, a spurt of invalid data will be sent to the vocoder which will cause a period of noise when the vocoder begins synthesizing. Also, extraneous data may cause the LPC algorithms to resynchronize to no avail. It becomes desirable to leave canned data in the buffers during the transition in order to produce silence from the new vocoder until valid data are received. A data pattern of alternating 0's and 1's will "silence" a continuously variable slope delta modulation (CVSD) algorithm, and a stream of 0's will "silence" an LPC algorithm and prevent resynchronization. However, either of these patterns will result in a fixed symbol which will prevent the particular modem simulated in this system from establishing synchronization during the start-up procedure. Bit patterns have been chosen to satisfy both the modem and LPC or CVSD upon switching rates. The 16-bit patterns are 0101 1010 0101 1010 for CVSD and 0001 0010 0100 1000 for LPC. No bit in the repeating pattern sent to either LPC will behave like a synchronization bit. The synchronization algorithm used by the LPCs has been improved to prevent erroneous resynchronization. The small buffer maintained by the CVSD algorithm in the LDSP is initially alternating 0's and 1's. The corresponding buffer for the LPC algorithm is initially all 0's, which prevents resynchronization; in addition, output to the D/A converter is forced to silence until valid data arrive which will cause the LPC to synchronize after twenty frames and start "speaking." In this way, the annoying noise associated with an out-of-synch LPC vocoder is eliminated.

A fixed probability of packet error on the nonexistent return link may be specified for any given demonstration of the system. The effect of any "errors" on the return link will be an increased delay in responding to changes in channel capacity. The complete system is operational

TABLE II PROBABILITY OF CORRECT SYMBOL RECEPTION				
Channel Capacity (bps)	Transmission Rate (bps)			
	2,560	5,120	10,240	20,480
40,960	1	1	1	1
20,480	1	1	1	1
10,240	1	1	1	0.99900
6,826	1	1	0.99994	0.98500
5,120	1	1	0.99900	0.94700
4,096	1	1	0.99480	0.89200
3,413	1	0.99994	0.98500	0.83000
2,925	1	0.99970	0.96900	0.76900
2,560	1	0.99900	0.94700	0.71200

with simulated rather than actual modem programs in LDVT-II and LDVT-III. In the simulated modem the multiple-frequency shift-keying (MFSK) channel signals are not created and decoded, but data are corrupted symbol by symbol (symbol = four bits) according to channel capacity and transmission rate, using the probabilities shown in Table II. No major surprises are anticipated when the final modem software is incorporated in the system.

III. ADAPTIVE VOICE COMMUNICATIONS IN A WIDEBAND MULTI-USER ENVIRONMENT

A. Simulation of a Single-Voice User in a Multi-user Environment

While the above-described real-time simulation allows us to assess the subjective effects of speaking over an adaptive time-variable link, real-time computational limitations restrict the bandwidth of that link to that of a single user. Wideband multi-user links (satellite channels for example) can introduce additional effects on received speech data that the single-user link simulation will not duplicate. Accordingly, we have begun to implement a hybrid (real-time plus non-real-time) facility that will allow a live speaker to communicate over a simulated wideband adaptive communications link. In brief, a non-real-time simulation will provide a history of packet loss, bit error, and voice coding bit rate for a single selected user in the wideband system. These parameters will be effected based on channel capacity changes and rate adaptation dynamics as well as by the activities of the other users of the wideband channel. The recorded history will later be imposed on live speech in a real-time simulation. The latter will consist of an LDSP for real-time voice analysis/synthesis and the PDP-11 as overall system controller. The PDP-11 will, at appropriate times, cause voice packets to be lost, bit errors to be introduced, or vocoder algorithms to switch in the LDSP, based on the recorded history produced by the earlier non-real-time simulation run. We will thus be able to experience live voice communications in a real-time environment and for arbitrary wideband link conditions and adaptation mechanics.

Efforts in this area resulted initially in a real-time PDP-11 program that interacts with a vocoder in the LDSP and discards voice packets in accordance with previously recorded measurements. The program distinguishes between silence and nonsilence intervals in the voice stream and only discards voice packets during actual talkspurts. This feature was included in order to more realistically simulate wideband links for which TASI-like protocols are employed. The program currently operates with LPC in the LDSP, and handles "discarded" packets by repeating the LPC parameters of the last received frame but with reduced excitation amplitude. Packet loss can be dictated either via an external data base produced by non-real-time simulation or based on a fixed probability specified at execution time. An existing non-real-time simulation of a fixed vocoder rate PVC voice network was modified to provide the packet loss history for a single user. This record was then applied to the real-time simulation to validate the system software.

Although the experiment just described involved fixed-rate speech, the real-time portion of this system allows rapid changes in vocoder bit rate. One of four vocoder algorithms operates in the LDSP depending on the current channel capacity. The capacity can be varied with time either in accordance with a preprogrammed sequence or in response to keyboard input. The rate of vocoder data transfer is governed by the A/D clock in the LDSP. The vocoders which currently interact with this half-duplex system are 2400- and 4800-bps LPC algorithms and 9600- and 19,200-bps CVSD algorithms. The four speech algorithms are stored in the PDP-11/45, and the vocoder switching delays increase in an almost linear manner as vocoder bit rate decreases. Changes in channel capacity and vocoder rate are indicated on a PDP-11/45 terminal. This real-time, variable-rate speech system has been made flexible enough so that a history of packet loss, bit errors, and coding rate gathered from a non-real-time facility could be imposed in a straightforward manner on live speech.

B. Analysis of Ground Station Parameters and Costs for Wideband Satellite Communications Experiments

1. Planned Broadcast Satellite Experiment

The real-time adaptive speech communication demonstration described in Sec. II simulates a situation where a variable-capacity line-of-sight (LOS) link is utilized for transmission of a single digitized speech signal. The link capacity directly determines what speech bit rate can be transmitted, and changes in this allowable bit rate are caused only by changes in channel noise level. The situation becomes more complicated when the communication medium is a wideband satellite channel shared among a number of ground stations and serving many voice and data users. Here the capacity available for a particular speech stream will depend not only on total satellite capacity but also on the fluctuating demands of other users.

An experimental network is being planned in which the issues of adaptive speech communication in a broadcast satellite environment can be investigated. The network initially will include four ground stations and a shared broadcast satellite channel of about 1.6-mbps capacity. These ground stations eventually will interface with a larger terrestrial network.

The purpose of this section is to discuss some of the considerations involved in obtaining satellite ground stations for this experiment. Of particular interest are the costs of various alternatives, and the advantages and disadvantages of operation in different frequency bands.

2. Satellite Link Relationships

In order to discuss ground station requirements, it is necessary to formulate the fundamental satellite link equations. The basic formula for the ratio of carrier power to noise density at a receiver with effective antenna aperture (efficiency times area) $\epsilon_R a_R$ and system temperature t_R, at a distance r from a transmitter with power p_T and antenna gain g_T, is written as

$$\left(\frac{c}{n_o}\right) = \frac{p_T g_T \epsilon_R a_R}{4\pi r^2 k t_R} \quad . \tag{1}$$

The gain aperture relation is

$$g = \frac{4\pi f^2}{c_o^2} \epsilon a \tag{2}$$

where c_o is the speed of light. The effective isotropic radiated power of a transmitting source is defined as

$$\mathrm{eirp} = p_T g_T \quad . \tag{3}$$

From Eqs. (1) and (2), the (c/n_o) received at the satellite can be expressed as

$$\left(\frac{c}{n_o}\right)_u = \frac{p_G \epsilon_G a_G g_S}{4\pi r^2 k t_S} \tag{4a}$$

where subscripts G and S refer to ground and satellite parameters, respectively. In decibels, this becomes

$$\left(\frac{C}{N_o}\right)_u = P_G + E_G + A_G - L + \left(\frac{G}{T}\right)_S + 228.6 \tag{4b}$$

where P_G is expressed in decibels with respect to one watt, A_G in decibels with respect to one square meter, and G/T for the satellite in decibels per degrees Kelvin, and where $E_G = 10 \log \epsilon_G$ and the spatial factor $L = 10 \log 4\pi r^2$. (Typically $E_G \approx -2.5$ dB for standard dishes at a wide range of frequencies.) If the full satellite eirp were transmitted as carrier, (c/n_o) at the ground receiver would be

$$\left(\frac{c}{n_o}\right)_d = \frac{(\mathrm{eirp})_S}{4\pi r^2} \frac{\epsilon_G a_G}{k t_G} \quad , \tag{5a}$$

or in decibels,

$$\left(\frac{C}{N_o}\right)_d = (\mathrm{EIRP})_S - L + E_G + A_G - T_G + 228.6 \tag{5b}$$

with T_G in degrees Kelvin. The carrier-to-noise-density ratio for the combined up-down link is

$$\left(\frac{c}{n_o}\right)_{ud} = \frac{1}{1/\left(\frac{c}{n_o}\right)_u + 1/\left(\frac{c}{n_o}\right)_d} \tag{6a}$$

Efforts in this area resulted initially in a real-time PDP-11 program that interacts with a vocoder in the LDSP and discards voice packets in accordance with previously recorded measurements. The program distinguishes between silence and nonsilence intervals in the voice stream and only discards voice packets during actual talkspurts. This feature was included in order to more realistically simulate wideband links for which TASI-like protocols are employed. The program currently operates with LPC in the LDSP, and handles "discarded" packets by repeating the LPC parameters of the last received frame but with reduced excitation amplitude. Packet loss can be dictated either via an external data base produced by non-real-time simulation or based on a fixed probability specified at execution time. An existing non-real-time simulation of a fixed vocoder rate PVC voice network was modified to provide the packet loss history for a single user. This record was then applied to the real-time simulation to validate the system software.

Although the experiment just described involved fixed-rate speech, the real-time portion of this system allows rapid changes in vocoder bit rate. One of four vocoder algorithms operates in the LDSP depending on the current channel capacity. The capacity can be varied with time either in accordance with a preprogrammed sequence or in response to keyboard input. The rate of vocoder data transfer is governed by the A/D clock in the LDSP. The vocoders which currently interact with this half-duplex system are 2400- and 4800-bps LPC algorithms and 9600- and 19,200-bps CVSD algorithms. The four speech algorithms are stored in the PDP-11/45, and the vocoder switching delays increase in an almost linear manner as vocoder bit rate decreases. Changes in channel capacity and vocoder rate are indicated on a PDP-11/45 terminal. This real-time, variable-rate speech system has been made flexible enough so that a history of packet loss, bit errors, and coding rate gathered from a non-real-time facility could be imposed in a straightforward manner on live speech.

B. Analysis of Ground Station Parameters and Costs for Wideband Satellite Communications Experiments

1. Planned Broadcast Satellite Experiment

The real-time adaptive speech communication demonstration described in Sec. II simulates a situation where a variable-capacity line-of-sight (LOS) link is utilized for transmission of a single digitized speech signal. The link capacity directly determines what speech bit rate can be transmitted, and changes in this allowable bit rate are caused only by changes in channel noise level. The situation becomes more complicated when the communication medium is a wideband satellite channel shared among a number of ground stations and serving many voice and data users. Here the capacity available for a particular speech stream will depend not only on total satellite capacity but also on the fluctuating demands of other users.

An experimental network is being planned in which the issues of adaptive speech communication in a broadcast satellite environment can be investigated. The network initially will include four ground stations and a shared broadcast satellite channel of about 1.6-mbps capacity. These ground stations eventually will interface with a larger terrestrial network.

The purpose of this section is to discuss some of the considerations involved in obtaining satellite ground stations for this experiment. Of particular interest are the costs of various alternatives, and the advantages and disadvantages of operation in different frequency bands.

or in decibels,

$$\left(\frac{C}{N_o}\right)_{ud} = 10 \log\left(\frac{c}{n_o}\right)_{ud} \quad . \tag{6b}$$

For a particular $(C/N_o)_{ud}$, the achievable bit rate is determined as

$$10 \log R = \left(\frac{C}{N_o}\right)_{ud} - \left(\frac{E_b}{N_o} + M\right) \tag{7}$$

where R is rate in bits per second, E_b is transmitted energy per bit in decibel-joules, and M is link margin. Achievement of a bit probability of error $<10^{-5}$ with standard modulation schemes requires E_b/N_o = 10 dB, and a conservative margin would be M = 10 dB.

3. Costs of Various Alternatives

Three major variable-cost items in the ground station are the antenna, the high-power transmitter amplifier (HPA), and the low-noise receiver amplifier (LNA). These components determine the parameters A_G, P_G, and T_G, respectively, in the link Eq.(4b) and (5b). The parameters $(EIRP)_S$ and $(G/T)_S$ relate to the satellite itself. A typical C-band (4/6 GHz) satellite such as Westar has an EIRP of 33 dBW and a G/T of −7 dB/°K. The proposed Satellite Business Systems (SBS) K_u-band (12/14 GHz) satellite will have an EIRP varying from 38 dBW in the Western U.S. to 42 dBW in the East, and a G/T varying from −6 dB/°K in the West to −2 dB/°K in the East. Taking the worst case (Western U.S.) figures for the SBS satellite, the gains in going to the higher frequency satellite are +5 dB in $(C/N_o)_d$ and +1 dB in $(C/N_o)_u$. The other elements, E_G and L in the link Eq.(4b) and (5b), can be taken as constants. The loss factor L = 163 dBm^2 at a typical range R = 24,600 mi, and in the calculations to follow the antenna efficiency loss E_G is assumed accounted for in the margin M, and therefore E_G is taken to equal zero.

In a recent report,* price information has been collected on antennas, HPAs, and LNAs at C- and K_u-bands from a number of vendors. A sampling of these cost data is summarized in Table III. Independent discussions with a vendor indicate that the K_u-band costs (note that costs are tabulated in thousand-dollar units) in this table may be slight underestimates, since they seem to be based on projected development of a substantial market for the K_u-band equipment. The same discussions indicated that costs at 18/30 GHz would be quite similar to costs at 12/14 GHz.

Application of Table III to determine achievable bit rate and cost for a particular alternative is illustrated as follows. Assume C-band operation with an antenna diameter d_G = 10 m, LNA noise temperature 150°K (corresponding to $t_G \approx 200$°K when 50°K for antenna noise is included), and transmitter power p_G = 20 W. In uplink Eq. (4b), P_G = 13 dBW, E_G = 0 (accounted for in M), $A_G = 10 \log (25\pi) = 19$ dBm^2, L = 163 dBm^2, and $(G/T)_S$ = −7 dB/°K (for Westar). The result is $(C/N_o)_u$ = 90.6 dB. In Eq.(5b), $(EIRP)_S$ = 33 dBW (for Westar) and $T_G = 10 \log 200 = 23$ dB/°K, so that $(C/N_o)_d$ = 94.6 dB. From Eq.(6b), the overall carrier-to-noise density is $(C/N_o)_{ud}$ = 89.1 dB. The achievable bit rate is obtained from Eq. (7) as 10 log R = 89.1 − 20 = 69.1, or R = 8.1 mbps.

* "Implications of Demand Assignment for Future Satellite Communication Systems," Final Report for Defense Communications Agency under Contract DCA100-76-C-0060, Systems Control, Inc. (June 1977).

TABLE III COSTS OF LNAs, HPAs, AND ANTENNAS		
	C-Band	K_u-Band
Noise Temperature t_{LNA} (°K)	LNA Cost ($K)	
110	5	35
150	5	25
300	3	14
Power P_G (W)	HPA Cost ($K)	
20	9	10
100	12	27
1000	50	140
Antenna Diameter d_G (m)	Antenna Cost ($K)	
2	2.5	3.5
3.2	5	5
4.5	12	12
10	40	–

TABLE IV

BIT RATES AND COSTS FOR ALTERNATIVE GROUND STATION CONFIGURATIONS

Frequency Band	d_G (m)	P_G (W)	t_{LNA} (°K)	R (mbps)	Cost of LNA + HPA + Antenna ($K)	Total Cost ($K)
C	10	20	150	8.1	54	138
C	4.5	20	150	1.6	26	100
K_u	4.5	20	300	2.3	36	110
K_u	3.2	100	150	4.0	57	131

This more than satisfies the 1.6-mbps requirement. From Table III, the (LNA + HPA + antenna) cost for this system is 5 + 9 + 40 = $54K. Cost could be reduced significantly (while still achieving the needed bit rate) by using a smaller antenna, but the Federal Communications Commission (FCC) currently requires a minimum of 9-m transmitter antenna at C-band to avoid domestic interference problems.

Table IV summarizes the results of rate and cost calculations for a variety of alternatives. The "total cost" column includes up- and down-converters ($12K each), miscellaneous equipment such as rack and dehydrator ($10K), system engineering and documentation ($15K), and site preparation and installation ($35K for 10-m antenna, $25K for smaller antenna). Modem cost is not included. The Western region SBS numbers given above were assumed for $(EIRP)_S$ and $(G/T)_S$ at K_u-band. The table indicates that the cost advantage of smaller antenna size at K_u-band is compensated for by the higher cost of 12/14-GHz transmitters and receivers. The major advantage of going to higher frequency appears to be the elimination of the domestic satellite interference problem (and the need for a waiver of the FCC minimum antenna size regulation) for antenna diameters less than 9 m. The difficulty with the higher frequencies is that commercial K_u-band satellites are not yet in operation, although SBS plans to begin operations in 1979.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

18 19

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: ESD-TR-77-296
2. GOVT ACCESSION NO.:
3. RECIPIENT'S CATALOG NUMBER: 7
4. TITLE (and Subtitle): 6 Information Processing Techniques Program. Volume II. Communications-Adaptive Internetting.
5. TYPE OF REPORT & PERIOD COVERED: Semiannual Technical Summary n, t. 1 Apr[redacted] - 3[redacted] Sep[redacted] [redacted]977
6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s): 10 Bernard Gold
8. CONTRACT OR GRANT NUMBER(s): 15 F19628-76-C-[redacted]2
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Lincoln Laboratory, M.I.T., P.O. Box 73, Lexington, MA 02173
10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: ARPA Order 2929, Program Element No. 62708E, 16 Project No. 7T10
11. CONTROLLING OFFICE NAME AND ADDRESS: Defense Advanced Research Projects Agency, 1400 Wilson Boulevard, Arlington, VA 22209
12. REPORT DATE: 11 3[redacted] Sep[redacted] [redacted]977
13. NUMBER OF PAGES: 18
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): Electronic Systems Division, Hanscom AFB, Bedford, MA 01731 12 21p.
15. SECURITY CLASS. (of this report): Unclassified

15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report): Approved for public release; distribution unlimited.
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)
18. SUPPLEMENTARY NOTES: Volume I is ESD-TR-77-248.
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

- information processing techniques
- communications-adaptive internetting
- digital voice communications
- packet speech
- network speech terminal
- C^2 links
- time-varying communications

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes work performed on the Communications-Adaptive Internetting program sponsored by the Information Processing Techniques Office of the Defense Advanced Research Projects Agency during the period 1 April through 30 September 1977.

DD FORM 1473, 1 JAN 73 — EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

207 650 4B